

Release Notes for SUSE Linux Enterprise Server 10 Service Pack 2

Release Notes for SUSE Linux Enterprise Server 10 Service Pack 2

Version 10.2.1, 13. May 2008

Abstract

These release notes are generic for all SUSE Linux Enterprise Server 10 based products. Some parts may not apply to a particular architecture/product. Where this is not obvious, the respective architectures are listed explicitly. The instructions for installing this Service Pack can be found in the `README` file on CD1. There are also translations of this file.

A startup and preparation guide are found under the `docu` directory on the media. Any documentation (if installed) can be found below `/usr/share/doc/` in the installed system.

This Novell product includes materials licensed to Novell under the GNU General Public License (GPL). The GPL requires that Novell make available certain source code that corresponds to those GPL-licensed materials. The source code is available for download at <http://www.novell.com/linux/source>. Also, for up to three years from Novell's distribution of the Novell product, upon request Novell will mail a copy of the source code. Requests should be sent by e-mail to sle_source_request@novell.com or as otherwise instructed at <http://www.novell.com/linux/source>. Novell may charge a fee to recover its reasonable costs of distribution.

1. Purpose	1
2. New Features	2
3. Driver Updates	6
3.1. Network Drivers	6
3.2. Storage Drivers	6
3.3. Other Drivers	7
4. Other Updates	8
5. Installation-Related Notes	10
6. Update-Related Notes	12
6.1. General Notes	12
6.2. Update from SUSE Linux Enterprise Server 10 SP1	12
6.3. Update from SUSE LINUX Enterprise Server 9	13
7. Technology Previews	16
8. Deprecated Functionality	18
9. Known Issues	19
10. Resolved Issues	22
11. Technical Informations	23
12. More Information and Feedback	25

Chapter 1. Purpose

This SUSE Linux Enterprise Server 10 Service Pack 2 serves several purposes:

- Provide enhancements to the SLES 10 code base (see Chapter 2, *New Features*).
- Provide all maintenance fixes (see Chapter 3, *Driver Updates*) released since GA of SLES 10.
- Provide an easy update (see README) of your system or individual packages to the latest Service Pack level. This is especially useful if you cannot use online update mechanisms.
- Provide an easy fresh install (see README) using the latest kernel, drivers, and installer updates.
- Include PTFs (special fixes for customers) which were folded back into the SLES 10 common code base making them part of the maintained code base.
- Provide useful additional information and documentation (see Chapter 5, *Installation-Related Notes*).

Through joint testing and maximum care, we try hard not to break any ISV certification with a Service Pack, but we recommend checking with your ISV about your application's certification status.

With the release of SUSE Linux Enterprise Server 10 Service Pack 2, the now obsoleted Service Pack 1 enters limited support status for the following 6 months, during which time Novell will continue to provide security updates and L3 support to maintain its customer's operations safe during the migration window. At the end of the six-month parallel support period, on November, 15th 2008, support for Service Pack 1 will be permanently discontinued.

Chapter 2. New Features

- 915resolution was added to change the resolution of the 845G, 855GM, 865G, 915G, 915GM, 945G, 945GM, 965G and 965GM chipsets
- Docking station support for Santa Rosa chipsets were added
- Complete iSCSI support was added to autoyast2
- Add support for --dynamic-list, AMDFAM10h, SSE5 and all POWER6 instructions to binutils
- Updated crash to version 4.0-4.8 and increased the number of supported CPUs to 16384
- Added support of dhcpFailOverPeer objects to dhcp
- Allow online resize of root filesystem on Ext3
- Enable krb4 support in evolution
- eDirectory NMAS support was added to FreeRADIUS
- Added optional SSE5 instructions to gcc
- Full RFC3484 (getaddrinfo/IPv6) support, AF_IUCV and support for up to 4096 CPUs added to glibc
- Added 'start this script before' support to insserv
- Added kdump package and YaST2 frontend with support for saving dumps over network and to filter kdump images.
- When doing a network install, check if the IP address of the machine to be installed is already in use by another machine. This check can be suppressed with parameter "forceip=1" either on the command line or in the configuration file.
- Added support for LUKS and LVM over iSCSI to mkinitrd
- New PAM module 'pam_faildelay' was added to change the delay on failure per application
- Added sbvim-cim-client, a CIM Client Class Library for Java applications
- Added StrongSwan 4.1.10, an OpenSource IPsec-based VPN Solution for Linux
- Added support for SGS Thomson Microelectronics Fingerprint Reader including a YaST2 module for configuration
- Added audit support to hwclock
- iSNS support was added to the YaST2 iSCSI modules
- Support for bridge devices, ficon devices and VLANs was added to the YaST2 network module
- The default disk label for disks larger than 2 TB is now GPT
- It is now possible to install on NFS as root partition
- The YaST2 storage module is now able to online resize ext3
- Python bindings for YaST2

- The maximum number of raw devices can be specified as boot option (`max_raw_minors=XX`, where $1 \leq XX \leq 65536$)
- Add a new kernel which is designed to run on top of a virtual machine interface layer (VMI). The VMI interface layer is part of the package `kernel-vmi` which is available as x86 package only (not for x86_64). This package can be installed via your favorite update tool once the system has a registered update channel.
- A hardware handler for the DM-MPIO framework to enable explicit ALUA support (ALUA: T10 Asymmetric LUN Access or T10 ALUA) was added.
- Support for Autofs version 5 was added additional to the standard version 4.
- The Open-iSCSI stack was updated to enhance the iSCSI boot/install capability so that it is not limited to a single integrated network interface any more. This update changed the `sysfs` layout.
- Support added for SB700 S/PATA
- Integrated support for AMDFam11h (Griffin)
- Integrated support for NFS locking of a clustered filesystem
- Support for mount counting and last check timestamping for `fsck -a` on `reiserfs` was added.
- AppArmor was updated to the latest version.
- Wide UTF-8 character support was enabled for Aspell.
- The following packages were added for better support of XEN management tools: `gtk-vnc`, `virt-viewer`
- The HPI STONITH module of heartbeat was enabled.
- Support for Intel(R) iAMT was added.
- IPA Fonts were added.
- The Firmware for Intel Wireless WiFi Link 4965AGN Cards was added.
- Added Unicode text editor 'yudit' to edit and convert text of different languages.
- Improved support for LVM and Device-Mapper based mirroring.
- SuSEfirewall2 now autodetect bridge interfaces and permit traffic.
- zSeries: AF_IUCV Protocol Support: Enables IUCV communication via the BSD socket interface between Linux VM guests or between a Linux guest and CMS. More info: Device Drivers, Features, and Commands - SC33-8289-04, Chapter 21. AF_IUCV protocol support.
- zSeries: Kernel NSS support: Saves memory when running multiple Linux server with the same kernel under z/VM by allowing to save the kernel in a Named Saved Segment in z/VM. More info: Device Drivers, Features, and Commands - SC33-8289-04, Chapter 13. Shared kernel support.
- zSeries: ETR (external time reference) Support: Enables Linux images to synchronize with parallel Sysplex or GDPS by providing z/OS compatible external timer reference and maintaining data consistency groups for XRC data mover. More info: The ETR support introduces a new kernel parameter "etr" that is used to set the initial state for the online attribute of the two ports. The syntax of the parameter is "etr=[on | off | port0 | port1]". The default is "off". The ETR support introduces a number of `sysfs`-

attributes. There are two time synchronization ports `etr0` and `etr1`, which can be accessed via sysfs under: `/sys/devices/system/etr`

- zSeries: Dynamic CHPID reconfiguration via SCLP: Allows to react to hardware changes and to configure path for the LPAR from the Linux system. More info: Device Drivers, Features, and Commands - SC33-8289-04, Chapter 2, Configuring a CHPID on LPAR and Chapter 32, `chcp` and `lschp`.
- zSeries: skb scatter-gather support for large incoming messages: Improves network throughput and reliability (avoid memory allocation failures in network subsystem) with large MTU sizes and jumbo frames. In particular this renders HiperSockets highly available even for large messages.
- zSeries: FCP performance data collection - adapter statistics: Provides more hardware performance metrics which are available from the FCP adapter on IBM System z9 and z10.
- zSeries: Support for Processor Degradation: Generates uevents for all CPUs if the CPU-capability changes. This can happen e.g. because the CPUs are overheating. The CPU-capability can be read via: `/sys/devices/system/cpu/cpuN/capability`. More info: Device Drivers, Features, and Commands - SC33-8289-04, Chapter 2, CPU capability change.
- zSeries: Provide Linux process data into z/VM monitor stream: Allows to consolidate monitor data from multiple z/VM guest with z/VM monitor tools. More info: Device Drivers, Features, and Commands - SC33-8289-04, Chapters 4, 14, 15, 16 and 32 (`mon_fsstatd` and `mon_procd`).
- zSeries: z/VM unit record device driver: Allows to access the z/VM reader from a running Linux system, useful for tasks like preparing the z/VM reader from Linux to install or to dump to the z/VM reader. More info: Device Drivers, Features, and Commands - SC33-8289-04, Chapter 18. z/VM unit record device driver.
- zSeries: STSI change for capacity provisioning: Make the permanent and temporary capacity information as provided by the STSI instruction of the IBM System z10 available to user space via `/proc/sysinfo`.
- zSeries: Hotplug daemon for CPU and memory management: New daemon in `s390-tools` 1.6.3 that manages CPU- and memory-resources based on a set of rules. Depending on the workload CPUs can be enabled or disabled. The amount of memory can be increased or decreased exploiting the Cooperative Memory Management (CMM1) feature. See: `man -a cpuplugd`. More info: Device Drivers, Features, and Commands - SC33-8411-00, Chapters 36 (`cpuplugd` - Activate CPUs and control memory Use).
- zSeries: HiperSockets MAC layer routing support: Enables Layer-2 support in HiperSockets, including IPv6 for Layer-2, on the new z10 hardware. Note that connecting Layer-2 and Layer-3 hosts is not supported. To use it, echo 1 > `layer2` attribute before setting the HiperSockets device online.
- zSeries: Support two OSA ports per CHPID: Exploits the next OSA adapter generation offering two ports within one CHPID, giving you four ports per card ("Four-port exploitation on OSA-Express3 GbE SX and LX"). The additional port number 1 can be specified with the `qeth` sysfs-attribute "`portno`". More info: Device Drivers, Features, and Commands - SC33-8411-00, Chapters 9 (Specifying the relative port number).
- zSeries: In kernel and software Support for CP Assist Instructions AES & SHA: Enables the crypto driver and library to exploit the new z10 hardware encryption providing performance enhancements for AES-192, AES-256, SHA-384 and SHA-512.
- zSeries: Linux CPU Node Affinity: Improve performance of the system by scheduling processes to the optimal node where the CPU is associated, exploiting the new z10 CPU node topology. More info: Device Drivers, Features, and Commands - SC33-8411-00, Chapters 2 (Examining the CPU topology).

- zSeries: Large Page Support: Enables better system performance with large memory footprints like in Java or database workloads by exploiting the new z10 hardware feature to be able to use large memory pages (1 MB). See: Linux kernel source tree under Documentation/vm/hugetlbpage.txt. More info: Device Drivers, Features, and Commands - SC33-8411-00, Chapters 28.

Chapter 3. Driver Updates

3.1. Network Drivers

- Added ipw3945 driver in version 1.2.0 and ipw3945d tools in version 1.7.22
- Added iwlwifi driver
- Added mac80211 wireless driver
- Updated Realtek r8169 driver to support newer chipsets
- Updated IPv6 stack for better Source Address Selection
- Updated tg3 driver to version 3.86b
- Added bnx2x 10 gigabit ethernet driver in version 1.42.3
- Updated bnx2 driver to version 1.6.7c
- Added ixgbe driver in version 1.1.21
- Updated e1000 driver to version 7.6.9.1
- Updated netxen driver to the one from upstream kernel version 2.6.23
- Updated qla3xxx driver to version 2.03.00-k4
- Added igb driver in version 1.0.8
- Added Chelsio 10GB driver (cxgb) in version 1.0
- Added Myricom 10G driver (myri10ge) in version 1.3.1
- Updated s2io driver to version 2.0.25.1

3.2. Storage Drivers

- Updated qla2xxx to version 8.02.00-k6
- Updated qla4xxx to version 5.01.02-d4
- Updated megaraid driver to version 2.20.5.2
- Updated megaraid_sas to version 3.15
- Updated 3w-9xxx to version 2.26.08.003 (9.5.0.1)
- Updated 3w-xxxx to version 1.26.02.002
- Updated aic7xxx to version 7.0
- Updated aic94xx driver to include SATA support for SAS
- Updated aacraid driver to version 1.1-5[2449]

- Updated cciss driver to version 3.6.20
- Updated lpfc driver to version 8.2.0.22
- Updated fusion driver to version 3.04.06
- Enable support for Nvidia MCP73, 77, 79 and 7B SATA devices

3.3. Other Drivers

- Updated CIFS to version 1.50c
- Updated OCFS2 to version 1.4.0
- Updated intel-i810 driver
- Added X11 driver for AMD Geode LX 2D (xorg-x11-driver-video-amd)
- Updated X11 driver for Radeon cards
- Updated XFS and DMAPI driver
- Updated XFS and DMAPI driver
- Updated Wacom driver to version 1.46

Chapter 4. Other Updates

- Updated heartbeat to version 2.1.3
- Updated ocfs2-tools to version 1.4.0
- Updated bin86 tools to version 0.16.17
- Updated dhcp6 to version 1.0.4
- Updated dmapi to version 2.2.8
- Updated heartbeat to version 2.1.3, including a resource agent for iSCSI initiator and hooks to SAP resource agents
- Updated iprutils to version 2.2.8
- Updated irqbalance to version 0.55 for better NUMA awareness
- Updated iscsitarget to version 0.4.15
- Updated k3b to version 1.0.4
- Updated ksh to version 0.93s+, changed signal handler for SIGINT and SIGQUIT to be compatible with the behavior on other operating systems and shells
- Updated libica to version 1.3.8
- Updated libvirt to version 0.3.3 including mDNS and NUMA support
- Updated lvspd to version 1.4.0
- Updated microcode_ctl to version 20080108
- Updated mod_mono to version 1.2.5
- Updated NetworkManager to version 0.6.5
- Updated nss_ldap to version 259
- Updated openCryptoki to version 2.2.4.1
- Updated openhpi to version 2.10.2
- Updated PHP to version 5.2.5
- Updated QT to version 4.3.4 for LSB compliance
- Updated scrim to version 1.4.7
- Updated sysstat to version 8.0.4
- Updated virt-manager to version 0.5.3
- Updated XEN to version 3.2
- Updated xfsdump to version 2.2.46

- Updated xfsprogs to version 2.9.4
- Updated xntp to version 4.2.4p3
- Updated xsp to version 1.2.5

Chapter 5. Installation-Related Notes

This section includes installation-related information for this Service Pack.

- Installation using Persistent Device names

If you plan to add additional storage devices to your system after the OS installation, we strongly recommend to use persistent device names for all storage devices during installation. The installer by default uses the kernel device names.

How to proceed:

During installation, enter the partitioner. For each partition, select "Edit" and go to the "FStab Options" dialog. Any mount option except "Device name" provide you persistent devicenames.

To switch an already installed system to using persistent device names, proceed as described above for all existing partitions. In addition, rerun the boot loader module in YaST to switch the bootloader to using the persistent device name also. Just start the module and select "Finish" to write the new proposed configuration to disk. This needs to be done before adding new storage devices.

For further information please look at http://en.opensuse.org/Persistent_Storage_Device_Names.

- MD Devices on top of iSCSI not possible

iSCSI devices cannot be used for Linux Software RAID. Using MD devices on top of iSCSI triggers a cyclic dependency that leads to a crash.

- Using qla3xxx and qla4xxx driver at the same time

QLogic iSCSI Expansion Card for IBM BladeCenter provides both Ethernet and iSCSI functions. Some parts on the card are shared by both functions. The current qla3xxx (ethernet) and qla4xxx (iSCSI) drivers support Ethernet and iSCSI function individually. They do not support using both functions at the same time. Using both Ethernet and iSCSI functions at the same time may hang the device and cause data lost and filesystem corruptions on iSCSI devices or network disruptions on Ethernet.

- Using iSCSI Disks When Installing

To use iSCSI disks during installation it is necessary to add the following parameter to the kernel parameter line: `withiscsi=1`

During installation, an additional screen appears that provides the possibility to attach iSCSI disks to the system and use them in the installation process.

Since SUSE Linux Enterprise Server 10 SP1 booting from an iSCSI server on i386, x86_64 and ppc is supported, when an iSCSI enabled firmware is used.

On ppc, a single bootfile (zImage.initrd) instead of yaboot is used.

- Using EDD Information for Storage Device Identification

If you want to use EDD information (`/sys/firmware/edd/<device>`) to identify your storage devices, change the installer default settings using an additional kernel parameter.

Requirements:

- BIOS provides full EDD information (found in `/sys/firmware/edd/<device>`)

- Disks are signed with a unique MBR signature (found in `/sys/firmware/edd/<device>/mbr_signature`)

Procedure:

- Add parameter `use_edd=1` to the kernel parameters during initial installation.
- The device-id list in the installer shows the EDD ID (such as `edd_dev80_part1`) instead of the default device-id name.
- Select those device IDs for installation and runtime (for example, in `/etc/fstab` and `bootloader`).
- Automatic installation with Autoyast in an LPAR (zSeries)

For automatic installation with Autoyast in an LPAR, it is required, that the parmfile used for such an installation has blank characters at the beginning and the end of each line (the first line need not to start with a blank). The number of character in one line should not exceed 80 characters.

- Adding DASD or zFCP disks during installation (zSeries)

The adding of DASD or zFCP disks is not only possible during the installation workflow, but also when the installation proposal is shown. To add disks at that stage please click on the "Expert" tab and scroll down. There the DASD and/or zFCP entry is shown. These added disks are not shown in the partitioner. To get the disks into the partitioner, you have to click on the expert label and select "reread partition table". This may reset any previously entered information.

- Creating LVM or EVMS Volumes with DASDs (zSeries)

If want to create a LVM or EVMS volume with DASDs that are not formatted or partitioned this will fail. The DASDs can be formatted in the DASD activation panel. Creating a partition can be done in the partitioner by hitting the [create] button and specifying "do not format" and removing any mountpoints indicated.

Chapter 6. Update-Related Notes

This section includes update-related information for this Service Pack.

6.1. General Notes

- SPident reports an old Service Pack level

SPident is a tool to identify the Service Pack level of the current installation. It may report that the system has not reached the level of this Service Pack. This happens, when *optional* updates that are not automatically installed by YOU are not manually selected during update. If you use or need any packages which have optional updates, select these in order to reach the current Service Pack level.

- Novell AppArmor

This release of SUSE Linux Enterprise Server ships with Novell AppArmor. The AppArmor intrusion prevention framework builds a firewall around your applications by limiting the access to files, directories, and POSIX capabilities to the minimum required for normal operation. AppArmor protection can be enabled via the AppArmor control panel, located in YaST under Novell AppArmor. For detailed information about using Novell AppArmor, see the documentation in `/usr/share/doc/packages/apparmor-docs`.

The AppArmor profiles included with SUSE Linux have been developed with our best efforts to reproduce how most users use their software. The profiles provided work unmodified for many users, but some users find our profiles too restrictive for their environments.

If you discover that some of your applications do not function as you expected, you may need to use the AppArmor Update Profile Wizard in YaST (or use the `aa-logprof(8)` command line utility) to update your AppArmor profiles. Place all your profiles into learning mode with the following: **`aa-complain /etc/apparmor.d/*`**

When a program generates many complaints, the system's performance is degraded. To mitigate this, we recommend periodically running the Update Profile Wizard (or `aa-logprof(8)`) to update your profiles even if you choose to leave them in learning mode. This reduces the number of learning events logged to disk, which improves the performance of the system.

- LD_ASSUME_KERNEL Environment Variable

Do not set the `LD_ASSUME_KERNEL` environment variable any longer. In the past, it was possible to use it to enforce LinuxThreads support, which was dropped. If you set `LD_ASSUME_KERNEL` to a kernel version lower than 2.6.5, everything breaks because `ld.so` looks for libraries in a version that does not exist anymore.

6.2. Update from SUSE Linux Enterprise Server 10 SP1

- New on disk format of new Sysstat package

The new features of the new Sysstat package needs a new on disk format of the data files. After the update of the sysstat package the old collected data can no longer be used.

- Changed order of starting network interface

The order in which network interfaces will be started has changed. The new order is now bond interfaces first, then vlan, dialup tunnel and finally bridge interfaces.

6.3. Update from SUSE LINUX Enterprise Server 9

- Becoming Superuser Using **su**

By default, calling **su** to become root does not set the PATH for root. Either call **su -** to start a login shell with the complete environment for root or set **ALWAYS_SET_PATH** to **yes** in **/etc/default/su** if you want to change the default behavior of **su**.

- Forwarding xauth keys between users with **sux**

The shell script **sux** was removed. The functionality of forwarding xauth keys between users is now handled by the **pam_xauth** module and **su**.

- NTP-Related Files Renamed

For reasons of compatibility with LSB (Linux Standard Base), most configuration files and the init script were renamed from **xntp** to **ntp**.

- Changed tar behavior in SUSE Linux Enterprise Server 10

Under SUSE Linux Enterprise Server 9, when extracting a directory from a tar archive that already existed as a symbolic link in the target directory, tar would overwrite the symlink with an actual directory. Under SUSE Linux Enterprise Server 10, tar leaves the symlink and places the contents of the archive within it.

To enforce the old behavior please use the option **--no-overwrite-dir** when extracting an archive.

- ulimits

SUSE Linux Enterprise Server 9 set up the user environment with an unlimited stack size resource limit to work around restrictions in stack handling of multithreaded applications. With SUSE Linux Enterprise Server 10, this is no longer necessary and has been removed. The login environment now defaults to the kernel default stack size limit. To restore the old behavior, add

```
ulimit -Ss unlimited
```

to **/etc/profile.local**. If you want an automatic configuration of your resource limits suited to protect desktop systems, you may want to install the **ulimit** package.

- Mounting Encrypted Partitions

With SUSE Linux Enterprise Server 10, we switched to "cryptoloop" as the default encryption module. SUSE Linux Enterprise Server 9 used twofish256 using **loop_fish2** with 256 bits. Now we are using twofish256 using **cryptoloop** with 256 bits. The old twofish256 is available as **twofishSL92**.

- Reconfiguring Intel and Nvidia Sound Drivers

When updating a system with the **snd-intel8x0** module (for Intel, SIS, AMD, and Nvidia on-board chips), the system might be unable to load the module at reboot, because the module option **joystick** was removed from the newer version. To fix the problem, reconfigure the sound system using YaST.

- Upgrading MySQL from SLES9 to SLES10

During the upgrade from SUSE Linux Enterprise Server9 to SUSE Linux Enterprise Server10 also MySQL is upgraded from 4.x to 5.x. To complete this migration you have also to upgrade your data as described in the MySQL documentation.

- Migrating from PHP 4 to PHP 5

Although most existing PHP 4 code should work without changes, there are a few backwards-incompatible changes. Find a list of these changes at: <http://www.zend.com/manual/migration5.incompatible.php>

- Switching from Heimdal to MIT Kerberos

MIT Kerberos is now used instead of heimdal. Converting an existing Heimdal configuration automatically is not always possible. During a system update, backup copies of configuration files are created in `/etc` with the suffix `.heimdal`. YaST-generated configuration settings in `/etc/krb5.conf` are converted, but check whether the results match your expectations.

Before starting the update, you should decrypt an existing Heimdal database into a human-readable file with the command

```
kadmin -l dump -d heimdal-db.txt
```

. This way, you can create a list of available principals that you can restore one-by-one using **kdc** from MIT Kerberos. Find more information about setting up a KDC in the documentation in the "krb5-doc" package.

To configure a Kerberos client, start the YaST Kerberos Client module and enter your values for "Standard Domain", "Standard Realm", and "KDC Server Address".

- MDNS and .local domain names

The .local top level domain is treated as link-local domain by the resolver. DNS requests are send as multicast DNS requests instead of normal DNS requests. If you already use the .local domain in your nameserver configuration you will have to switch this option off in `/etc/host.conf`. Please also read the `host.conf` manual page, more information on multicast DNS can be found on <http://www.multicastdns.org>.

MDNS can be disabled during installation by booting with the `nomdns` option set.

- Fine-Tuning Firewall Settings

SuSEfirewall2 is enabled by default. That means that by default you cannot log in from remote systems. It also interferes with network browsing and multicast applications, such as SLP and Samba ("Network Neighborhood"). You can fine-tune the firewall settings using YaST.

- CD/DVD device name on pSeries changed

With SUSE Linux Enterprise Server 10 SP1, the built-in CD/DVD drive on POWER3/POWER4 pSeries models p610/p615/p630 will be accessed with the libata kernel driver because it is more reliable. On all POWER5 models the libata driver is used to allow DLPAR hotplug operations.

This changes the kernel device name from `/dev/hda` to `/dev/sr0`.

- vsftpd with xinetd

Starting with SUSE Linux Enterprise Server 10, vsftpd can be configured independently or over the xinetd. The default is stand-alone. In previous versions, the default was xinetd.

To run it over xinetd, make sure that the service is enabled in the xinetd configuration (`/etc/xinetd.d/vsftpd`) and set the following line in `/etc/vsftpd.conf`:

```
listen=NO
```

- Setting Up D-BUS for Interprocess Communication in `.xinitrc`

Many applications now rely on D-BUS for interprocess communication (IPC). Calling **dbus-launch** starts **dbus-daemon**. The systemwide `/etc/X11/xinit/xinitrc` uses `dbus-launch` to start the window manager.

If you have a local `~/.xinitrc` file, you must change it accordingly. Otherwise applications might fail. Save your old `~/.xinitrc`. Then copy the new template file into your home directory with:

```
cp /etc/skel/.xinitrc.template ~/.xinitrc
```

Finally, add your customizations from the saved `~/.xinitrc`.

- Modular KDB

KDB is no longer available as a loadable module on all architectures except Itanium. KDB is only supported in the debug kernel.

- PCMCIA

cardmgr no longer manages PC cards. Instead, as with Cardbus cards and other subsystems, a kernel module manages them. All necessary actions are executed by **hotplug**. The `pcmcia` start script has been removed and **cardctl** is replaced by **pccardctl**. For more information, see `/usr/share/doc/packages/pcmciautils/README.SUSE`.

Chapter 7. Technology Previews

Technology Preview features are not supported or only supported limitedly. These features are mainly included for customer convenience and may not be functionally complete, instable or in other ways not suitable for production use.

- Hot-Add of Memory

Hot-Add-memory is currently only supported on the following machines:

- IBM xSeries x260
- IBM xSeries single node x460
- IBM xSeries x3800
- IBM xSeries x3850
- IBM xSeries single node x3950

If your machine is not listed, please call support, whether the machine has been successfully tested. Else a maintenance update will explicitly mention the general availability of this feature.

- Oracle Cluster File System on zSeries

Due to known issues, OCFS2 is not recommended for production use on zSeries (not supported). A maintenance update will explicitly mention the availability of this function.

- Huge Page Memory support via HMC on POWER

Huge Page Memory support (16GB pages, enabled via HMC) is not yet supported under Linux. Problems occur if huge pages are assigned to a partition in combination with eHEA / eHCA adapters. eHEA: Network interfaces can't be setup if huge page memory is assigned to the same partition.

- libhugetlbfs

The libhugetlbfs project shipped with SUSE Linux Enterprise Server 10 is a preview of application provision with transparent access to system huge pages. While the library provides an application with easy access to huge pages when sufficient huge pages have been previously allocated on the system, additional development and testing is required to provide a stable transition to normal pages in a production environment.

- Read-Only Root Filesystem

It is possible to run SUSE Linux Enterprise Server 10 from Service Pack 2 on on a read-only root filesystem. Due to the huge number of possible configurations, this is currently not a supported scenario.

The `/tmp` and `/var/tmp` directories needs to be on a separate partition and cannot be mounted read-only.

After the installation has finished and all services are configured, login as `root` and do the following modifications:

Modify `/etc/fstab` and add "ro" to the mount options of the root filesystem entry.

```
rm /etc/mtab
ln -s /proc/mounts /etc/mtab
mkdir /var/lib/hwclock
mv /etc/adjtime /var/lib/hwclock
ln -s /var/lib/hwclock/adjtime /etc/adjtime
# the following two steps are only necessary if you use dhcp:
mv /etc/resolv.conf /var/lib/misc/
ln -s /var/lib/misc/resolv.conf /etc/resolv.conf
# Now mount root filesystem read-only and reboot
mount -o remount,ro /
reboot
```

Chapter 8. Deprecated Functionality

The following list of current functionality is deprecated and will be removed with the next Service Pack or major SUSE Linux Enterprise Server release.

- The JFS filesystem is no longer supported for new installations. The kernel file system driver is still there, but YaST does not offer partitioning with JFS.
- For the future strategy and development with respect to volume- and storage-management on SUSE Linux Enterprise System, please see: <http://www.novell.com/linux/volumemanagement/strategy.html>
- The **ippl** package is deprecated and will be removed with SUSE Linux Enterprise Server 11.
- **powertweak** package is deprecated and will be removed with SUSE Linux Enterprise Server 11.
- CTC, ESCON, and IUCV IP interfaces are no longer officially supported. For compatibility reasons, they are still usable, but with the next release of SUSE Linux Enterprise Server, the support of these interfaces will be dropped completely.
- For reasons of compatibility with SUSE Linux Enterprise Server 9, the mapped-base functionality is present in SUSE Linux Enterprise Server 10. This functionality is used by 32-Bit applications that need a larger dynamic data space (such as database management systems).

With SUSE Linux Enterprise Server 10, a similar functionality called flexmap is available. Because this is now the preferred way, mapped-base is deprecated and will vanish in future releases.

Chapter 9. Known Issues

- i586 and i686 Machine with more than 16 GB of Memory

Depending on the workload, i586 and i686 machines with 16GB-48GB of memory can run into instabilities. Machines with more than 48GB of memory are not supported at all. To run on such a machine, lower the memory with the `mem=` kernel boot option.

On such memory scenarios we strongly recommend to use a x86-64 with 64-bit SUSE Linux Enterprise Server and run the x86 applications on it.

- Bootloader and mount by UUID or LABEL

When the way the root device is mounted (by UUID or by label) is changed in YaST, the boot loader configuration needs to be saved again to make the change effective for the boot loader.

The "mount by" setting displayed in the YaST2 boot loader module is the setting that will be in effect after saving the configuration.

- EVMS Volumes Might Not Appear When Using iSCSI

If you have installed and configured an iSCSI SAN and have created and configured EVMS Disks or Volumes on that iSCSI SAN, your EVMS volumes might not be visible or accessible. This problem is caused by EVMS starting before the iSCSI service. iSCSI must be started and running before any disks or volumes on the iSCSI SAN can be accessed.

To resolve this problem, enter either **`chkconfig evms on`** or **`chkconfig boot.evms on`** at the Linux server console of every server that is part of your iSCSI SAN. This ensures that EVMS and iSCSI start in the proper order each time your servers reboot.

- cpio and files larger 4GB

cpio is not able to add files larger than 4GB to an archive.

- KDE and IPv6 Support

By default, IPv6 support is not enabled for KDE. You can enable it using the `/etc/sysconfig` editor of YaST. This feature is disabled because IPv6 addresses are not properly supported by all Internet service providers and, as a consequence, would lead to error messages while browsing the Web and delays while displaying Web pages.

- Installing/Updating on IBM System z9

When installing SUSE Linux Enterprise Server 10 on a System z9, some restrictions apply through hardware or software. Some of these restrictions are part of these Release Notes. For an updated list, refer to http://www-128.ibm.com/developerworks/linux/linux390/october2005_restrictions.html.

For IBM System z9 machines ensure to have MCF RJ9967101E or IBM System z9 GA3 base driver installed. Otherwise Linux reboot will not work.

- Using Disks in z/VM (zSeries)

If SUSE Linux Enterprise Server 10 is installed on disks in z/VM, which reside on the same physical disk, the created access path (`/dev/disk/by-id/`) is not unique. The ID of a disk is the ID of the underlying disk. So if two or more disk are on the same physical disk, they all have the same ID.

To avoid this ambiguity, please use the access path by-path. This can be specified during the installation when the mount points are defined.

The above restriction does not apply for SLES10 SP2 (which has fix for Problem-ID 34345 and 43704), if you have the z/VM PTF for APAR VM64273, with which z/VM provides a unique identifier that allows to distinguish between virtual disks on the same real device. Udev rules that provide both old and new /dev/disk/by-id paths are included. To use the new IDs in your multipath setup, please replace the `getuid_callout "/sbin/dasinfo -u -b %n"` with `"/sbin/dasinfo -x -b %n"` in your multipath configuration for DASD devices. Please see the man page for the `dasinfo` tool for additional information.

- Local Mounts of iSCSI Shares

An iSCSI shared device should never be mounted directly on the local machine. In an OCFS2 environment, doing so causes all hardware to hard hang.

- Restriction When Using `cpint/hcp` (zSeries)

When using the `cpint/hcp` interface with z/VM 5.1 or earlier, the guest should not have more than 2 GByte of storage. If the guest has more storage, the command may fail.

- HP sx1000- and sx2000-based systems and multiple VGA adapters

HP sx1000- and sx2000-based systems, including rx7620, rx8620, rx7640, rx8640 and all Superdomes, only support the use of a single VGA device per hard partition. The usage of more than one VGA card causes a system hang or crash on all of this systems.

To prevent this on machines with more than one VGA adapter, either do not install the X server at all or configure it manually to use only a single VGA device.

- YaST2 CD-Creator and YUM installation sources

The YaST2 CD-Creator module does not support YUM installation sources like our update server provides. For this reason, it is not possible to create a medium with updates included. If you want to create a medium with updates included, use YaST2 Product-Creator.

YaST2 Product-Creator is a successor of YaST2 CD-Creator. It includes a GUI for the kiwi imaging system. This way it is also possible to create a Live-CD, XEN image, etc. from the same configuration used in the CD Creator. The Product-Creator will get shipped together with the SDK.

- `glibc-debuginfo` package on POWER

The `glibc-debuginfo` package doesn't set up the debug information files for certain processor-optimized versions of the libraries. This affects `glibc` debugging on the following POWER processors: POWER4, PPC970, POWER5+, POWER6x. To fix the problem, one needs to manually create the missing symbolic links using a sequence of commands like this:

Obtain root privileges:

```
$ /bin/su -
```

Go to the 32-bit debuginfo directory:

```
# cd /usr/lib/debug/lib
```



```
# ln -s ppc970 power4
# ln -s power5 power5+
# ln -s power6 power6x
```

Go to the 64-bit debuginfo directory:

```
# cd /usr/lib/debug/lib64
# ln -s ppc970 power4
# ln -s power5 power5+
# ln -s power6 power6x
# ln -s . tls
```

Chapter 10. Resolved Issues

- Bugfixes

This Service Pack contains all the latest bugfixes for each package released via the maintenance Web since the GA version.

- Security Fixes

This Service Pack contains all the latest security fixes for each package released via the maintenance Web since the GA version.

- Program Temporary Fixes

This Service Pack contains all the PTFs (Program Temporary Fix) for each package released via the maintenance Web since the GA version which were suitable for integration into the maintained common codebase.

Chapter 11. Technical Informations

This section contains a number of technical changes and enhancements for the experienced user.

- Locale Settings in `~/ .i18n`

If you are not satisfied with locale system defaults, change the settings in `~/ .i18n`. Entries in `~/ .i18n` override system defaults from `/etc/sysconfig/language`. Use the same variable names but without the `RC_` namespace prefixes, for example, use `LANG` instead of `RC_LANG`. For information about locales in general, see "Language and Country-Specific Settings" in the Reference Manual.

- Configuration of `kdump`

The kernel is crashing or otherwise misbehaving and a kernel core dump needs to be captured for analysis.

A description on how to setup `kdump` can be found under the following URL: http://www.novell.com/support/search.do?cmd=displayKC&docType=kc&externalId=3374462&sliceId=SAL_Public

- Realtime Applications

When running real-time applications on larger systems, lower maximum latencies can be achieved by employing the new `disable_buffer_lru` kernel command-line option. This disables the per-CPU LRU in the buffer cache, and may thus decrease overall filesystem performance.

- JPackage Standard for Java Packages

Java packages are changed to follow the JPackage Standard (<http://www.jpackage.org/>). Read the documentation in `/usr/share/doc/packages/jpackage-utils/` for information.

- Loading unsupported kernel drivers

To load unsupported kernel drivers automatically during boot, set the `sysconfig` variable `LOAD_UNSUPPORTED_MODULES_AUTOMATICALLY` in `/etc/sysconfig/hardware/config` to "yes".

- Nonexecutable Stack

Already introduced for SUSE Linux Enterprise Server 9 on the x86-64 (AMD64) architecture with 64-bit kernels, the Linux kernel in SUSE Linux Enterprise Server also supports nonexecutable stack (NX) on x86 for CPUs that support it (Intel Prescott and AMD64) with 32-bit kernels. For this to work, the kernel with PAE support, `kernel-bigsmp`, must be installed. Go into YaST and install that kernel instead of your default kernel. For 64-bit kernels, all kernels support NX.

The nonexecutable stack improves the security of your system. Many security vulnerabilities are stack overflows, where an attacker overwrites the stack of your program by feeding oversized data to the application that fails to properly check the length. Depending on the details of the program, with nonexecutable stack, these vulnerabilities may either not be exploitable (and only crash the program, resulting in a DoS) or at least be significantly harder to exploit.

Some applications do require executable stacks. The compiler detects this during compilation and marks the binaries accordingly. The kernel enable an executable stack for them to allow them to work.

On x86-64, to provide a higher level of security, the user can pass `noexec=on` on the kernel command line. The kernel then uses a nonexecutable stack unconditionally and also marks the data section of a

program nonexecutable. This provides a higher protection level than just the nonexecutable stack, but potentially causes problems for some applications. Novell has not found any problems during testing the most commonly used applications and services. Because it is not the default, this has not been tested as extensively as the stack protection alone, so Novell only recommends this setup for servers after the administrator has verified that all needed services continue to function properly.

Chapter 12. More Information and Feedback

- Read the READMEs on the CDs.
- Get the detailed changelog information about a particular package from the RPM:

```
rpm --changelog -qp <FILENAME>.rpm
```

<FILENAME>. is the name of the RPM.

- Check the ChangeLog file in the top level of CD1 for a chronological log of all changes made to the updated packages.
- Find more information in the docu directory of CD1 of the SUSE Linux Enterprise Server 10 CDs. This directory includes PDF versions of the SUSE Linux Enterprise Server 10 startup and preparation guides.
- <http://www.novell.com/documentation/sles10/> contains additional or updated documentation for SUSE Linux Enterprise Server 10.
- Visit <http://www.novell.com/linux/> for the latest Linux product news from SUSE/Novell and <http://www.novell.com/linux/source/> for additional information on the source code of SUSE Linux Enterprise products.